arXiv:1510.04524v1 [cs.IT] 15 Oct 2015

Old Bands, New Tracks—Revisiting the Band Model for Robust Hypothesis Testing

Michael Fauß *Student Member, IEEE* and Abdelhak M. Zoubir, *Fellow, IEEE*

***Abstract*—The density band model proposed by Kassam [1] for robust hypothesis testing is revisited in this paper. First, a novel criterion for the general characterization of least favorable distributions is proposed, which unifies existing results. This criterion is then used to derive an implicit definition of the least favorable distributions under band uncertainties. In contrast to the existing solution, it only requires two scalar values to be determined and eliminates the need for case-by-case statements. Based on this definition, a generic fixed-point algorithm is proposed that iteratively calculates the least favorable distributions for arbitrary band specifications. Finally, three different types of robust tests that emerge from band models are discussed and a numerical example is presented to illustrate their potential use in practice.**

***Index Terms*—Robust hypothesis testing, robust detection, distributional robustness, model uncertainties, band model, mismatch.**

I. Introduction

STATISTICAL hypothesis tests are referred to as *robust*, if they are insensitive to small, random deviations from the underlying model. In this paper we consider robustness against *distributional uncertainties*, meaning that, under either hypothesis, the distribution of the observed random variable is only known approximately. Each hypothesis is hence composite, i.e., represented by a set or class of possible distributions. A test is further called *minimax* robust, if it guarantees a certain maximum error probability over the entire set of distributions specified by the composite hypotheses. Because of this property, minimax robust tests are often essential for the design of systems that have to function reliably in harsh environments or cannot be modeled accurately.

The field of robust statistics, and robust hypothesis testing in particular, was developed foremost by Huber in the mid-1960s [2], [3]. He was the first to derive the famous clipped likelihood ratio test, which is robust against outliers of the ε-contamination type, i.e., infrequent, grossly corrupted observations. This kind of contamination is particularly critical since a single corrupted observation can be enough to alter the outcome of a non-robust test [4]. The clipped likelihood ratio test was further shown to be a test of two simple hypotheses. More precisely, it is a regular likelihood ratio test of the so-called *least favorable* instead of the *nominal* distributions, the latter denoting the distributions of the uncontaminated data.

M. Fauß and A. M. Zoubir are with the Signal Processing Group, Institute for Telecommunications, Department of Electrical Engineering and Information Technology, Technische Universität Darmstadt, Merckstr. 25, 64283 Darmstadt, Germany. E-mail: {mfauss,zoubir}@spg.tu-darmstadt.de.

This work was performed within the LOEWE Priority Program Cocoon (www.cocoon.tu-darmstadt.de) supported by the LOEWE research initiative of the state of Hesse/Germany.

Manuscript submitted October 8, 2015

Despite their wide use in practice, ε-contamination models are often not sufficient to accurately describe the uncertainty in the distributions. On the one hand, the assumption that the majority of the data follows the nominal distribution exactly can be too optimistic. On the other hand, the assumption that the outliers are drawn from an arbitrary distribution can be too pessimistic. In addition, many types of uncertainties cannot easily be incorporated into an ε-contamination model. Approximately known shapes or positions of distributions, for example, are usually hard to formulate in terms of nominals and outliers. In such cases, techniques involving the estimation of the true model may be preferable, like adaptive nonlinearities [5] or generalized likelihood ratio tests [6]. However, in contrast to minimax solutions, such methods do not guarantee pre-specified error probabilities and require changes in the distributions to happen slowly enough to update the estimates. For these reasons, more flexible uncertainty models for minimax robust tests are a topic of ongoing research [7]–[9].

In 1981, Kassam published a paper titled "Robust Hypothesis Testing for Bounded Classes of Probability Densities" [1] in which he proposed an uncertainty model for probability distributions that later became known as the *band model*. It allows each hypothesis to be formulated in terms of a density band, within which the true density is supposed to lie, and generalizes the outlier models suggested by Huber [3], Österreicher [10], Levy [11] and others. It can further be interpreted as both an ε-contamination model with bounded outlier distributions, or a model for general uncertainties in the shape that can be specified without introducing nominals. A more detailed discussion is deferred to later sections.

Considering its generality and versatility, it is astonishing that the band model has received very little attention in the robust statistics literature. While Huber's seminal paper on ε-contaminated observations has enjoyed an unbroken stream of citations since its publication in 1965, Kassam's paper, although covering a more general case, did not have a comparable impact.

One of the reasons for the limited interest in [1] might be the form of its main result. The theorem stating the least favorable densities spans the space of a column and distinguishes between four special cases, each involving a piecewise definition of the densities. In order to know which case holds, one has to check the existence or non-existence of in total six constants that have to be chosen such that the solutions are valid densities. In some cases the solution involves a function that "can be found", but is not specified explicitly. Even though none of these issues is critical, the solution and its calculation appear to be cumbersome and inelegant, especially compared to the lean results for the ε-

contamination model. Many authors seem to have come to the conclusion that the additional complexity the band model introduces to the least favorable densities is too high a price to pay for the increase in generality.

As a consequence, the popularity of the band model in robust hypothesis testing has been steadily decreasing. While it used to be one of the standard models in robust testing and filtering in the 1980s [12], [13], today many signal processing practitioners are not aware of its existence and some books on robust statistics ignore it entirely [11], [14].

In this paper, we revisit the band model, give a detailed description of its properties and motivate its use in practice. Beyond that, our original contribution is threefold: First, we state a novel criterion for the characterization of least favorable distributions. In our view, this criterion is conceptually simpler than the existing ones and results in more tractable optimization problems. Second, we derive a concise implicit definition of the least favorable distributions for the band model that offers additional insight into their structure. Third, based on this definition, we propose a fixed-point algorithm that provides a simple generic alternative for calculating the densities of the least favorable distributions without the need for a case analysis, such as in [1].

The paper is organized as follows: In Section II, we give a brief overview of some fundamental concepts in minimax robust hypothesis testing. The question of how to characterize least favorable distributions in general is discussed in Section III. After introducing the band model in Section IV, we derive its least favorable distributions in Section V. In Section VI, a fixed-point algorithm for their calculation is stated. A detailed discussion of the different types of robust tests that result from band models, is given in Section VII. In Section VIII, we present an example for how the band model can be used in practice.

Notation: We denote random variables with upper case letters and their realizations with the corresponding lower case letters. Similarly, probability measures (distributions) are denoted by upper case letters, the corresponding densities by lower case letters. The notation $\{X = x\}$ is shorthand for $\{\omega \in \Omega : X(\omega) = x\}$ and $P[X = x]$ for $P(\{X = x\})$. The expected value of a random variable X with respect to a measure P is written as $E_P[X]$. Occasionally, a tuple (x_0, x_1) is referred to simply as x. This will be clear from the context.

II. Fundamentals of Minimax Robust Detection

Let $(X_1, \dots, X_n)$ be a sequence of independently and identically distributed random variables with common distribution P defined on some measurable space $(\Omega, \mathcal{F})$. Throughout the paper, we assume that all probability measures have a continuous density function with respect to some common measure μ, i.e.,

$$\int_B dP = \int_B p \, \mathrm{d}\mu, \quad \forall B \in \mathcal{F}.$$

We denote the set of all distributions on $(\Omega, \mathcal{F})$ that admit this property by $\mathcal{M}_\mu$.

The goal of a binary statistical hypothesis test is to decide between the two hypotheses

$$\begin{aligned} \mathcal{H}_0 : &\quad P = P_0 \\ \mathcal{H}_1 : &\quad P = P_1, \end{aligned}$$

where $P_0, P_1 \in \mathcal{M}_\mu$ are two given distributions and $\mathcal{H}_0$ and $\mathcal{H}_1$ are referred to as the null and alternative hypothesis, respectively. A statistical test for $\mathcal{H}_0$ against $\mathcal{H}_1$ is in general defined by a decision $d \in \{0, 1\}$ and a (randomized) decision rule

$$\delta : \Omega^n \to [0, 1],$$

where $\delta = \delta(x_1, \dots, x_n)$ denotes the conditional probability to decide for the alternative hypothesis, given the observations $(x_1, \dots, x_n)$. The set of all decision rules is denoted by Δ. The type I and type II error probabilities are given by

$$\begin{aligned} P_0[d = 1] &= E_{P_0}[\, \delta \,], \\ P_1[d = 0] &= E_{P_1}[1 - \delta]. \end{aligned}$$

The optimal decision rule δ^* for the simple binary hypothesis test is a threshold comparison of the likelihood ratio, i.e.,

$$\delta^* = \begin{cases} 1, & z_n > \eta \\ \kappa, & z_n = \eta \\ 0, & z_n < \eta, \end{cases} \tag{1}$$

where $\eta > 0$ is the threshold value, $\kappa \in [0, 1]$ can be chosen arbitrarily, and $z_n : \Omega^n \to \mathbb{R}_+$ denotes the likelihood ratio

$$z_n := \prod_{i=1}^{n} \frac{\mathrm{d}P_1}{\mathrm{d}P_0}(x_i) = \prod_{i=1}^{n} \frac{p_1(x_i)}{p_0(x_i)}.$$

The likelihood ratio test is optimal in a very general sense [15]. In particular, it minimizes the weighted sum error probability, i.e., it solves

$$\min_{\delta \in \Delta} \; E_{P_0}[\, \delta \,] + \lambda \, E_{P_1}[1 - \delta], \tag{2}$$

where the weighting factor $\lambda = 1/\eta \geq 0$ determines the threshold value in (1). The robust version of (2) is considered in the following sections. However, we want to emphasize that the resulting minimax solution is optimal also in a Neyman–Pearson and a Bayesian sense [3].

In robust testing, the distribution under each hypothesis is assumed not to be known exactly. This distributional uncertainty is modeled by two disjoint sets $\mathcal{P}_0, \mathcal{P}_1 \subset \mathcal{M}_\mu$ so that the hypotheses become

$$\begin{aligned} \mathcal{H}_0 : &\quad P \in \mathcal{P}_0, \\ \mathcal{H}_1 : &\quad P \in \mathcal{P}_1. \end{aligned}$$

Tests for two sets of distributions are known as composite hypothesis tests and have been studied extensively in the literature [15], [16]. What distinguishes minimax robust procedures from other approaches is that the test is designed *a priori* so as to guarantee a certain reliability for all possible pairs $(P_0, P_1) \in \mathcal{P}_0 \times \mathcal{P}_1$. Mathematically, this property is formulated in terms of a minimax problem. The robust testing problem corresponding to (2) is thus given by

$$\min_{\delta \in \Delta} \max_{(P_0, P_1) \in \mathcal{P}_0 \times \mathcal{P}_1} \; E_{P_0}[\, \delta \,] + \lambda \, E_{P_1}[1 - \delta]. \tag{3}$$

A decision rule and a pair of distributions that solve (3) are called minimax optimal. The existence and characterization of minimax optimal solutions, however, is an intricate question. It has received continuous attention in applied mathematics, physics, and economics since the 1950s and is still an active area of research [17]–[19].

The most useful minimax theorem in the context of robust hypothesis testing is due to Sion [20]. In a nutshell, it follows from Sion's minimax theorem that a minimax optimal test exists for compact convex sets $\mathcal{P}_0, \mathcal{P}_1$ and is again a likelihood ratio test of the form (1), but with the nominal distributions P_0 and P_1 replaced by the least favorable distributions Q_0 and Q_1. The design of minimax robust tests hence reduces to finding the pair (Q_0, Q_1). This, however, is a non-trivial problem in itself. In the next section, we review two commonly used characterizations for least favorable distributions and propose a third one that we find to be simpler and more suitable in an optimization context.

III. Criteria for Least Favorable Distributions

The first criterion for least favorable distributions was given by Huber [2], who used it to prove minimax optimality of the clipped likelihood ratio test. It is now known by the name *stochastic dominance*.

Theorem 1 (Stochastic Dominance): A pair of distributions $(Q_0, Q_1) \in \mathcal{P}_0 \times \mathcal{P}_1$ is least favorable for all sample sizes $N \geq 1$, i.e., minimax optimal in combination with a likelihood ratio test, if it fulfills

$$\begin{aligned} Q_0\left[\frac{q_1}{q_0} > \eta\right] &\geq P_0\left[\frac{q_1}{q_0} > \eta\right] \quad \text{and} \\ Q_1\left[\frac{q_1}{q_0} \leq \eta\right] &\geq P_1\left[\frac{q_1}{q_0} \leq \eta\right] \end{aligned} \tag{4}$$

for all $(P_0, P_1) \in \mathcal{P}_0 \times \mathcal{P}_1$ and all $\eta \geq 0$.

The interpretation of Theorem 1 is that the distributions Q_0 and Q_1 have to be chosen such that for a single sample test the probabilities of both error types are jointly maximized, irrespective of the value of the likelihood ratio threshold. Stochastic dominance is a natural and widely used criterion for minimax optimality [8], [21].

About a decade later, Huber and Strassen derived a different, more technical criterion to characterize least favorable distributions.

Theorem 2 (Theorem 6.1 in [22]): Let $\Psi : [0,1] \to \mathbb{R}$ be a twice continuously differentiable function. A pair of distributions $(Q_0, Q_1) \in \mathcal{P}_0 \times \mathcal{P}_1$ is least favorable in the sense of Theorem 1, if it minimizes

$$H_\Psi(P_0, P_1) = \int \Psi\left(\frac{p_0}{p_0 + p_1}\right) \mathrm{d}(P_0 + P_1) \tag{5}$$

for all Ψ and among all $(P_0, P_1) \in \mathcal{P}_0 \times \mathcal{P}_1$.

Theorem 2 implies that the least favorable distributions concurrently minimize *all* f-divergences among $(P_0, P_1) \in \mathcal{P}_0 \times \mathcal{P}_1$.

Qualitatively speaking, the idea underlying stochastic dominance is that the least favorable distributions should maximize the error probabilities, while Theorem 2 corresponds to the more abstract intuition that the least favorable distributions should be "as similar as possible".

We propose a criterion that lends itself to both interpretations and, as we believe, helps to simplify and unify the theory of robust testing.

Theorem 3: A pair of distributions $(Q_0, Q_1) \in \mathcal{P}_0 \times \mathcal{P}_1$ is least favorable in the sense of Theorem 1, if it maximizes

$$L_\lambda(P_0, P_1) = \int \min\{p_0, \lambda\, p_1\} \,\mathrm{d}\mu \tag{6}$$

for all $\lambda \geq 0$ and among all $(P_0, P_1) \in \mathcal{P}_0 \times \mathcal{P}_1$.

A proof is detailed in Appendix A. The two aspects of error maximization and distance minimization can both be found in (6). For example, evaluating (2) for $N = 1$ yields

$$\begin{aligned} \min_{\delta \in \Delta} E_{P_0}[\delta] + \lambda E_{P_1}[1 - \delta] &= \min_{\delta \in \Delta} \int \delta p_0 + \lambda(1 - \delta) p_1 \,\mathrm{d}\mu \\ &= \int \min\{p_0, \lambda p_1\} \,\mathrm{d}\mu, \end{aligned}$$

which is the expression given in Theorem 3. This means that the least favorable distributions concurrently maximize the weighted error sum of a single sample test for all weighting coefficients $\lambda \geq 0$, which is in close analogy to the stochastic dominance characterization. At the same time, it can be shown that maximizing (6) is equivalent to maximizing

$$\int \min\left\{\frac{p_0}{p_0 + p_1}, \frac{\lambda}{1 + \lambda}\right\} \mathrm{d}(P_0 + P_1),$$

which, disregarding the differentiability assumption, is a special case of Theorem 2 with Ψ chosen to be $\Psi(s) = -\min\{s, \lambda/(1+\lambda)\}$.

IV. The Band Model

The band model proposed by Kassam in [1] covers composite hypotheses of the form

$$\mathcal{P}^{=} = \{P \in \mathcal{M}_\mu \,:\, p' \leq p \leq p''\}, \tag{7}$$

where p' and p'' fulfill

$$0 \leq p' \leq p'', \quad P'(\Omega) \leq 1, \quad P''(\Omega) \geq 1.$$

In words, it restricts the true density to lie within a band specified by p' and p''. This is indicated by the notation $\mathcal{P}^{=}$. Note that P' and P'' are measures on $(\Omega, \mathcal{F})$, but not probability measures, and P'' does not need to be finite.

Alternatively, the band model can be interpreted as an ε-contamination model with bounded outlier distribution. In this view, (7) can equivalently be written as

$$\mathcal{P}^{=} = \{P \in \mathcal{M}_\mu \,:\, P = P' + \varepsilon H \,,\, H \leq P'' - P'\}, \tag{8}$$

where P' corresponds to the scaled nominal distribution,

$$\varepsilon = 1 - P'(\Omega) \tag{9}$$

is the contamination rate and $H \in \mathcal{M}_\mu$ the outlier distribution, which is now bounded by $P'' - P'$. In this regard, the band model is an ε-contamination model that allows the incorporation of *a priori* knowledge in form of additional constraints on the outlier distribution.

V. Least Favorable Distributions for the Band Model

In this section, we state and discuss the main result of the paper, which is an implicit characterization of the pair of least favorable distributions for the band model (7).

A. An Implicit Characterization

Theorem 4: Given hypotheses of the form (7), the pair $(Q_0, Q_1) \in \mathcal{P}_0^{=} \times \mathcal{P}_1^{=}$ is least favorable in the sense of Theorem 1, if the densities (q_0, q_1) satisfy

$$\begin{aligned} q_0 &= \min\{p_0'' \,,\, \max\{c_0(\alpha q_0 + q_1) \,,\, p_0'\}\} \\ q_1 &= \min\{p_1'' \,,\, \max\{c_1(q_0 + \alpha q_1) \,,\, p_1'\}\} \end{aligned} \tag{10}$$

for some $\alpha \geq 0$ and some $c_0, c_1 \in (0, 1/\alpha]$. Such a pair always exists.

The proof of Theorem 4 is given by first deriving an upper bound on the function L_λ in Theorem 3 and then showing that the densities in (10) attain this bound. Their existence is shown in Section VI by means of a constructive algorithm.

Theorem 5: Given hypotheses of the form (7), for all $(P_0, P_1) \in \mathcal{P}_0^{=} \times \mathcal{P}_1^{=}$ it is the case that L_λ in Theorem 3 is upper bounded by

$$L_\lambda(P_0, P_1) \leq \int \min\{\hat{q}_0 \,,\, \lambda\, \hat{q}_1\}\, \mathrm{d}\mu + v_0\varepsilon_0 + \lambda v_1 \varepsilon_1,$$

where

$$\hat{q}_i = v_i p_i' + (1 - v_i) p_i'', \quad i = 0, 1,$$

$\varepsilon_0, \varepsilon_1 \in [0, 1]$ are defined in (9) and $v_0, v_1 \in [0, 1]$ can be chosen arbitrarily.

A proof is laid down in Appendix B. In the following paragraphs, it will become clear that the bound in Theorem 5 can be tightened to $v_0, v_1 \in \{0, 1\}$. We chose to state it in this more relaxed form to emphasize that an upper bound on L_λ can be obtained from every convex combination of the upper and lower bound.

We now show that for every feasible combination of λ, α and (c_0, c_1), the densities in (10) attain the upper bound in Theorem 5 for some choice of v_0 and v_1. Four cases are covered separately, but only two of them in detail since the others follow analogously.

Case 1: $c_0 \geq 1/(\alpha + \lambda)$ and $c_1 > \lambda/(1 + \alpha\lambda)$

On the set $\{q_1 > \lambda q_0\}$ it holds that

$$\begin{aligned} q_1 &> \lambda q_0 \\ \alpha q_0 + q_1 &\leq (\alpha + \lambda) q_0 \\ \frac{1}{\alpha + \lambda}(\alpha q_0 + q_1) &\leq q_0. \end{aligned} \tag{11}$$

Inserting q_0 from (10) into the right hand side of (11) yields

$$\frac{1}{\alpha + \lambda}(\alpha q_0 + q_1) > \max\{p_0'' \,,\, \max\{c_0(\alpha q_0 + q_1) \,,\, p_0'\}\},$$

which implies

$$q_0 = p_0'' \quad \text{for} \quad q_1 > \lambda q_0. \tag{12}$$

Analogously, on the set $\{q_1 \leq \lambda q_0\}$ it holds that

$$\begin{aligned} q_1 &\leq \lambda q_0 \\ (1 + \alpha\lambda) q_1 &\leq \lambda(q_0 + \alpha q_1) \\ q_1 &\leq \frac{\lambda}{1 + \alpha\lambda}(q_0 + \alpha q_1). \end{aligned} \tag{13}$$

Inserting q_1 from (10) into the left hand side of (13) yields

$$\min\{p_1'' \,,\, \max\{c_1(q_0 + \alpha q_1) \,,\, p_1'\}\} \leq \frac{\lambda}{1 + \alpha\lambda}(q_0 + \alpha q_1)$$

which can only be satisfied if

$$p_1'' \leq \frac{\lambda}{1 + \alpha\lambda}(q_0 + \alpha q_1).$$

This in turn implies

$$q_1 = p_1'' \quad \text{for} \quad q_1 \leq \lambda q_0. \tag{14}$$

Combining (12) and (14) yields

$$\int \min\{q_1 \,,\, \lambda q_0\}\, \mathrm{d}\mu = \int \min\{p_1'' \,,\, \lambda p_0''\}\, \mathrm{d}\mu,$$

which is the upper bound in Theorem 5 evaluated at $v_0 = v_1 = 0$.

Case 2: $c_0 < 1/(\alpha + \lambda)$ and $c_1 \leq \lambda/(1 + \alpha\lambda)$

On $\{q_1 \leq \lambda q_0\}$ we have

$$\frac{1}{\alpha + \lambda} q_1 \leq \min\{p_0'' \,,\, \max\{c_0(\alpha q_0 + q_1) \,,\, p_0'\}\}$$

and consequently

$$q_0 = p_0' \quad \text{for} \quad q_1 \leq \lambda q_0. \tag{15}$$

This is to say, that the least favorable distribution q_0 equals its lower bound on $\{q_1 \leq \lambda q_0\}$. Consequently, the outlier distribution $H_0 = Q_0 - P_0'$ is concentrated on $\{q_1 > \lambda q_0\}$. Analogously, on $\{q_1 > \lambda q_0\}$

$$\min\{p_1'' \,,\, \max\{c_1(q_0 + \alpha q_1) \,,\, p_1'\}\} > \frac{\lambda}{1 + \alpha\lambda}(q_0 + \alpha q_1),$$

holds, which implies that

$$q_1 = p_1' \quad \text{for} \quad q_1 > \lambda q_0 \tag{16}$$

so that H_1 is concentrated on $\{q_1 \leq \lambda q_0\}$. Combining (16) and (15) yields

$$\begin{aligned} \int \min\{q_1, \lambda q_0\}\, \mathrm{d}\mu &= \int \min\{p_1' + h_1, \lambda(p_0' + h_0)\}\, \mathrm{d}\mu \\ &= \int \min\{p_1', \lambda p_0'\}\, \mathrm{d}\mu + \varepsilon_1 + \lambda\varepsilon_0, \end{aligned}$$

which is the upper bound in Theorem 5 evaluated at $v_0 = v_1 = 1$.

Case 3: $c_0 \geq 1/(\alpha + \lambda)$ and $c_1 \leq \lambda/(1 + \alpha\lambda)$

Combining the arguments from the previous cases, it can easily be shown that

$$q_0 = p_0'' \quad \text{for} \quad q_1 > \lambda q_0,$$

and

$$q_1 = p_1' \quad \text{for} \quad q_1 \leq \lambda q_0,$$

which corresponds to $v_0 = 0$ and $v_1 = 1$ in Theorem 5.

Case 4: $c_0 < 1/(1 + \lambda)$ and $c_1 > \lambda/(1 + \lambda)c_0$

Again, using the same arguments as before, we obtain

$$q_0 = p_0' \quad \text{for} \quad q_1 > \lambda q_0,$$

and

$$q_1 = p_1'' \quad \text{for} \quad q_1 \leq \lambda q_0,$$

which corresponds to $v_0 = 1$ and $v_1 = 0$ in Theorem 5. □

B. Discussion

First, we point out two interesting special cases of Theorem 4. Choosing $\alpha = 0$ yields

$$\begin{aligned} q_0 &= \min\{p_0'' \,,\, \max\{c_0 q_1 \,,\, p_0'\}\}, \\ q_1 &= \min\{p_1'' \,,\, \max\{c_1 q_0 \,,\, p_1'\}\}. \end{aligned} \tag{17}$$

This is likely to be the most intuitive form of Theorem 4 and the most useful in practice—see Section VI. It closely resembles the structure of the asymptotic minimax solution $\tilde{Q}_0$, $\tilde{Q}_1$ derived in [23], which for general uncertainty sets $\mathcal{P}_0$ and $\mathcal{P}_1$ is given by

$$\begin{aligned} \tilde{Q}_0 &= \min_{P_0 \in \mathcal{P}_0} D_{\mathrm{KL}}(P_0 \| \tilde{Q}_1), \\ \tilde{Q}_1 &= \min_{P_1 \in \mathcal{P}_1} D_{\mathrm{KL}}(\tilde{Q}_0 \| P_1), \end{aligned} \tag{18}$$

where D_{KL} denotes the Kullback-Leibler (KL) divergence. In (18), $\tilde{Q}_0$ ($\tilde{Q}_1$) is the projection of $\tilde{Q}_1$ ($\tilde{Q}_0$) onto $\mathcal{P}_0$ ($\mathcal{P}_1$) with respect to the KL divergence. In (17) this projection is performed with respect to general f-divergences and the projection operator is written out explicitly. Note that asymptotically least favorable distributions exist for every convex uncertainty set, while strictly least favorable distributions do not.

A problem that arises when stating Theorem 4 with $\alpha = 0$ is that (17) is a sufficient but not a necessary condition for Q_0 and Q_1 to be least favorable. Assume, for example, that the two density bands do not overlap, i.e., $\{p_0'', p_1'' > 0\} = \emptyset$. In this case c_0 and c_1 can be chosen arbitrarily large without ever producing valid densities on the right hand side of (17). Nevertheless, every feasible pair is least favorable in the sense of Theorem 1. By adding a second density to the term that is scaled, this problem can be avoided. By choosing the least favorable distribution itself, it is guaranteed that optimality still holds.

In general, α can be chosen arbitrarily. However, a second noteworthy special case is $\alpha = 1$, for which Theorem 4 becomes

$$\begin{aligned} q_0 &= \min\{p_0'' \,,\, \max\{c_0(q_0 + q_1) \,,\, p_0'\}\}, \\ q_1 &= \min\{p_1'' \,,\, \max\{c_1(q_0 + q_1) \,,\, p_1'\}\}. \end{aligned}$$

In this form, Q_0 and Q_1 are the projections of the distribution $\frac{1}{2}(Q_0 + Q_1)$ onto $\mathcal{P}_0^=$ and $\mathcal{P}_1^=$, respectively. Since these projections are unique, knowledge of $\frac{1}{2}(Q_0 + Q_1)$ is sufficient to obtain Q_0 and Q_1. In this sense there exists a *single* least favorable distribution whose projections onto the respective bands form the least favorable *pair*. In applications, this property might be used to trade memory for computing power by storing only $\frac{1}{2}(Q_0 + Q_1)$ and calculating Q_0, Q_1 on demand.

We also point out that irrespective of the choice for α in Theorem 4, two constants c_0, c_1 are sufficient to characterize the solution. On the one hand, this is a major simplification compared to the six constants introduced in [1]. On the other hand, it is in perfect analogy to the ε-contamination model and shows how closely the two models are related.

Let us investigate this relation a bit further. The likelihood ratio of the densities in Theorem 4 can take on six possible values, namely

$$\frac{q_1}{q_0} \in \left\{ \frac{p_1'}{p_0'} \,,\, \frac{p_1''}{p_0'} \,,\, \frac{p_1'}{p_0''} \,,\, \frac{p_1''}{p_0''} \,,\, \frac{1 - \alpha c_0}{c_0} \,,\, \frac{c_1}{1 - \alpha c_1} \right\}.$$

Note that some of the terms can be zero or involve a division by zero. This corresponds to observations that are possible only under one of the hypotheses and hence lead to an unambiguous decision against the impossible hypothesis. Setting $p_0'' = p_1'' = \infty$ and $\alpha = 0$, we obtain Huber's least favorable densities for the ε-contamination model[1]

$$\begin{aligned} q_0 &= \max\{c_0 q_1 \,,\, p_0'\}, \\ q_1 &= \max\{c_1 q_0 \,,\, p_1'\}, \end{aligned}$$

with the corresponding clipped likelihood ratio

$$\frac{q_1}{q_0} \in \left\{ \frac{p_1'}{p_0'} \,,\, \frac{1}{c_0} \,,\, c_1 \right\}.$$

These basic derivations show how the ε-contamination model emerges naturally as a special case of the band model, which is in contrast to the prevailing perception that density bands are a rather tedious generalization of ε-contamination.

VI. Computation of the Least Favorable Densities

In this section we propose a fixed-point algorithm that makes use of the implicit definition in Theorem 4 to successively approximate q_0 and q_1. It offers a generic, conceptually simple and easy-to-implement way to determine the least favorable densities without the need for a case analysis.

First, let the functions $g_i : \mathbb{R} \to \mathbb{R}_+$ be defined by

$$g_i(c\,;\, p) := \int \min\{p_i'' \,,\, \max\{c\, p \,,\, p_i'\}\} \,\mathrm{d}\mu - 1$$

where $i = 0, 1$ and $p > 0$ is any μ-integrable function.

Lemma 1: Given any pair of distributions $P_0 \in \mathcal{P}_0^=$, $P_1 \in \mathcal{P}_1^=$, the functions $g_0(c_0\,;\, \alpha p_0 + p_1), g_1(c_1\,;\, p_0 + \alpha p_1)$ are non-decreasing and continuous and there exist some $c_0^*, c_1^* \in (0, \frac{1}{\alpha}]$ such that $g_0(c_0^*\,;\, \alpha p_0 + p_1) = 0$ and $g_1(c_1^*\,;\, p_0 + \alpha p_1) = 0$.

A proof is detailed in Appendix C. The algortihm we propose is given in Table VI. It reduces the problem of determining the least favorable densities to a repeated search for the root of a monotonic and continuous function. Convergence of Algorithm 1 is proven in Appendix D.

The least favorable densities q_0 and q_1 are not necessarily unique, cf. [1]. Therefore, the solution of Algorithm 1 depends on the initial densities q_0^0 and q_1^0. The minimax optimality of the robust test is not affected by this dependence.

The termination criterion in line seven is intentionally left vague. One option is to require that $\max_{i \in \{0,1\}} \| q_i^n - q_i^{n-1} \|$

[1]The case-by-case definition given by Huber is more common in the literature. It can be shown that both definitions are identical.

1: **input:** p_0', p_0'', p_1', p_1'', $\alpha \geq 0$
2: **initialize:** Choose feasible initial densities $q_0^0 \in \mathcal{P}_0^=$, $q_1^0 \in \mathcal{P}_1^=$ and set $n \leftarrow 0$.
3: **repeat**
4: Set $n \leftarrow n+1$
5: Solve $g_0(c_0\,;\alpha q_0^{n-1} + q_1^{n-1}) = 0$ for c_0 and set

$$q_0^n \leftarrow \min\{p_0''\,,\, \max\{c_0(\alpha q_0^{n-1} + q_1^{n-1})\,,\, p_0'\}\}.$$

6: Solve $g_1(c_1\,;q_0^n + \alpha q_1^{n-1}) = 0$ for c_1 and set

$$q_1^n \leftarrow \min\{p_1''\,,\, \max\{c_1(q_0^n + \alpha q_1^{n-1})\,,\, p_1'\}\}.$$

7: **until** $q_0^n \approx q_0^{n-1}$ and $q_1^n \approx q_1^{n-1}$
8: **return** q_0^n, q_1^n

TABLE I
ITERATIVE ALGORITHM TO CALCULATE THE LEAST FAVORABLE DENSITIES FOR THE BAND MODEL (7).

is sufficiently small, where $\|\cdot\|$ denotes a suitable norm. Alternatively, some f-divergence between Q_0^n and Q_1^n or the likelihood ratio can be tracked.

The speed of convergence depends on the choice of α. In general, α should be chosen as small as possible to achieve fast convergence and in many cases $\alpha = 0$ is the best option. Nonzero α-values are necessary only when the joint support of the least favorable distributions is very small. The easiest way to identify such cases is running Algorithm 1 with $\alpha = 0$ and checking whether or not a solution for c_0 and c_1 exists.

VII. Robust Tests Resulting from Band Models

Often, statements can be found in the literature which claim that "the robust solution for the band model [...] is similar to that for the ε-contamination model" [12]. In this section we show that this is not true in general. More precisely, we identify three characteristic types of robust tests that a band model can produce. These are explained and illustrated using a simple Gaussian example with lower bounds

$$\begin{aligned} p_0'(\omega) &= 0.8\, p_{\mathcal{N}}(\omega\,;-1,2),\\ p_1'(\omega) &= 0.8\, p_{\mathcal{N}}(\omega\,;1,2), \end{aligned} \tag{19}$$

where $p_{\mathcal{N}}(\omega\,;m,\sigma)$, $\omega \in \Omega = \mathbb{R}$ denotes the density of a Gaussian distribution with mean m and standard deviation σ. In terms of an outlier model, this corresponds to a 20% contamination rate and nominals

$$\begin{aligned} p_0(\omega) &= p_{\mathcal{N}}(\omega\,;-1,2),\\ p_1(\omega) &= p_{\mathcal{N}}(\omega\,;1,2). \end{aligned} \tag{20}$$

The three types of robust test are obtained by either *clipping*, *censoring* or *compressing* the nominal test statistic.

1) *Clipping:* For very large upper bounds p_0'', p_1'' the band model reduces to the ε-contamination model and the least favorable test becomes a clipped likelihood ratio test. This means that the influence of a single observation is bounded and seemingly very significant observations are trusted only to a certain extend. An example is shown in Fig. 1a, where the upper bounds have been chosen as ten times the nominals:

$$p_0'' = 10\, p_0 \quad \text{and} \quad p_1'' = 10\, p_1. \tag{21}$$

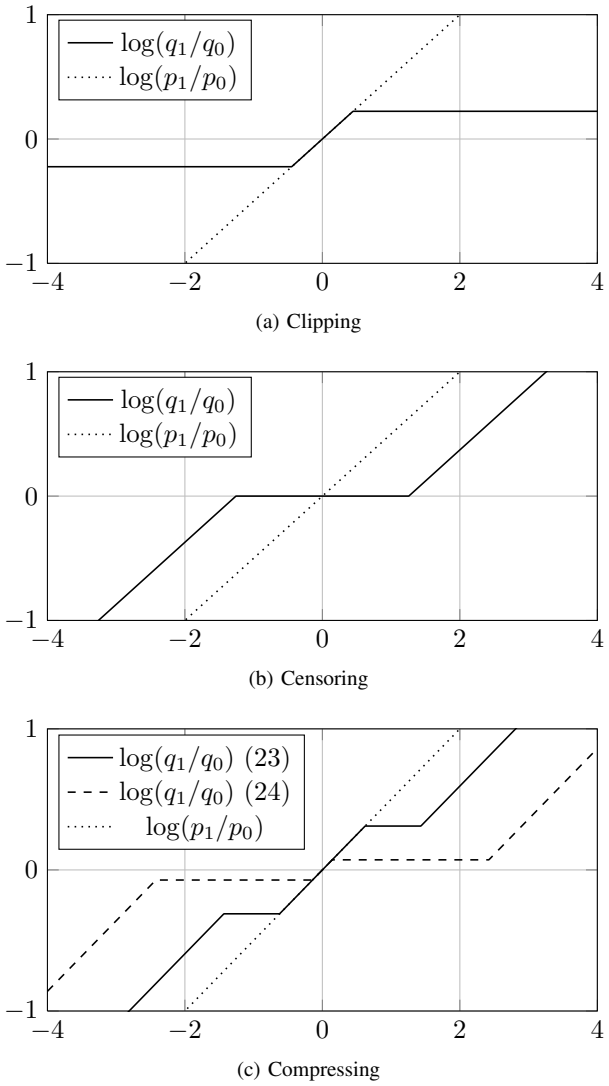

(a) Clipping

(b) Censoring

(c) Compressing

Fig. 1. Three different types of robust tests resulting from different band models. The least favorable densities are denoted by q_0, q_1, the nominals, which are given in (20), by p_0, p_1. The lower density bounds for all cases are given in (19). The upper bounds for the clipping case (a) are given in (21), for the censoring case (b) in 22) and for the compression case (c) in (23) and (24), respectively.

2) *Censoring:* For moderately chosen upper bounds the band model often results in likelihood ratios that are *censored* around one. An example with

$$p_0'' = 1.5\, p_0 \quad \text{and} \quad p_1'' = 1.5\, p_1. \tag{22}$$

is shown in Fig. 1b. This type of robust test can be seen as a kind of opposite of the clipped test. The influence of a highly significant observation is downweighted, but still unbounded. In contrast, observations whose significance is below a certain threshold are ignored entirely. This behavior of robust tests has been observed before in the context of *model uncertainties*, meaning that the data is not corrupted by gross outliers, but that the true distributions differ slightly from the nominal model. Examples for such uncertainty classes are the bounded KL divergence and bounded Hellinger distance classes, whose least favorable distributions have been shown to result in censored likelihood ratio tests [8], [11].

The intuition behind censored likelihood ratio tests is that under model uncertainties an observation needs to have a certain significance in order to reliably associate it with a hypothesis. Below this level, there is no clear preference and the observation should simply be ignored.

It is worth mentioning that censoring appears at likelihood ratios close to one, but not necessarily exactly one. In fact, the effectiveness of censoring depends on choosing the likelihood ratio threshold in accordance with the censoring level, otherwise the test is still minimax optimal with respect to the weighted sum error probability, but highly biased towards one hypothesis.

3) *Compressing:* So far, we have covered the cases of no uncertainty (nominal likelihood ratio test), moderate uncertainty (censored likelihood ratio test) and high uncertainty (clipped likelihood ratio test). In the transition regions between the nominal and the censored, as well as between the censored and the clipped test, the band model results in a likelihood ratio that has a plateau at some intermediate level. The examples in Fig. 1c illustrate this phenomenon. They have been obtained by choosing

$$p_0'' = 1.2\, p_0 \quad \text{and} \quad p_1'' = 1.2\, p_1 \tag{23}$$

and

$$p_0'' = 2.5\, p_0 \quad \text{and} \quad p_1'' = 2.5\, p_1 \tag{24}$$

respectively. It can be seen that the robust test statistic is the nominal likelihood ratio up to a certain threshold, a constant value on an interval of medium significance and a scaled version (shifted in the log-domain) of the nominal likelihood ratio on regions of high significance. The test statistics are neither clipped nor censored, but the influence of most observations is reduced—the more significant they are, the more pronounced the reduction. We hence refer to this as a *compressed* test statistic.

From the limited number of experiments we performed, it seems that compression is the most frequent outcome of a band model. In general, it offers a good tradeoff between test performance and robustness with respect to outliers as well as model uncertainties. A more realistic example discussed in the next section illustrates this.

VIII. Example

An idea that arises naturally when dealing with band models is to use confidence intervals of density estimators to construct the bands. Even though this approach has been suggested repeatedly in the literature [1], [12], we are not aware of any experiments. In this section, we present the results of a simple experiment that can be seen as a proof of concept for combining confidence intervals and band models.

We consider a problem that arises, for example, in spectrum sensing applications [24], such as cognitive radio [25], where a secondary user has to reliably detect ongoing transmissions of a primary user, in order to opportunistically occupy or free the channel. Since the secondary users are often equipped with battery-powered devices and the spectrum has to be sensed frequently, low-complexity energy detectors are a popular choice for this task.

The most commonly used signal model in this context is a complex Gaussian signal in complex Gaussian noise. The hypotheses are accordingly given by

$$\begin{aligned} \mathcal{H}_0 :&\quad X_n = W_n, \\ \mathcal{H}_1 :&\quad X_n = S_n + W_n, \end{aligned}$$

where $n = 1, \dots, N$ and W_n and S_n are independently distributed circular symmetric zero-mean complex Gaussian random variables with standard deviations σ_{W} and σ_{S}, respectively. The distribution of $|X_n|^2$ can be shown to be

$$p_0(x\,;\,\sigma_{\mathrm{W}}^2) = \frac{1}{\sigma_{\mathrm{W}}^2} e^{-\frac{x}{\sigma_{\mathrm{W}}^2}}, \quad x \in \mathbb{R}_+ \tag{25}$$

and

$$p_1(x\,;\,\sigma_{\mathrm{W}}^2, \sigma_{\mathrm{S}}^2) = \frac{1}{\sigma_{\mathrm{S}}^2 - \sigma_{\mathrm{W}}^2} \left(e^{-\frac{x}{\sigma_{\mathrm{S}}^2}} - e^{-\frac{x}{\sigma_{\mathrm{W}}^2}} \right) \tag{26}$$

under $\mathcal{H}_0$ and $\mathcal{H}_1$, respectively.

The noise and signal powers may fluctuate rapidly over time [26], making it difficult to estimate them reliably. For this example we assume that $\sigma_{\mathrm{W}}^2 \in [1, 2]$ and $\sigma_{\mathrm{S}}^2 \in [4, 6]$. It is not hard to show that in this case the least favorable distributions in the sense of Theorem 1 are

$$q_0(x) = p_0(x\,;\,2) = \tfrac{1}{2} e^{-\frac{x}{2}} \tag{27}$$

and

$$q_1(x) = p_1(x\,;\,2, 4) = \tfrac{1}{2} \left(e^{-\frac{x}{4}} - e^{-\frac{x}{2}} \right). \tag{28}$$

Obviously, this model is highly simplified. The purpose of introducing it here is to investigate how well the least favorable densities (27) and (28) can be reproduced without any knowledge of the underlying model, rather by means of training data and the band model. The procedure we follow is simple, but not unlike procedures used in practice:

1) Take N_0 (N_1) samples under hypothesis $\mathcal{H}_0$ ($\mathcal{H}_1$).
2) Calculate density estimates $\hat{p}_0$ and $\hat{p}_1$.
3) Calculate confidence intervals $\hat{p}_0', \hat{p}_0''$ and $\hat{p}_1', \hat{p}_1''$.
4) Calculate $\hat{q}_0, \hat{q}_1$ using the confidence intervals as density bands.

For the experiment, $N_0 = N_1 = 400$ samples were generated under each hypothesis, assuming that noise and signal power vary with every sample and are uniformly distributed over the respective interval. For the density estimation, we used the second of the two kernel density estimators detailed in [27]. It is tailored for densities on the nonnegative reals and uses gamma kernels. The bandwidth was determined via least-squares cross validation. The estimator was then applied to 500 data sets that were bootstraped from the original data [28]. The pointwise maximum and minimum of the corresponding density estimates were used as confidence intervals. The resulting density bands are shown in Fig. 2.

The least favorable densities for this band model were calculated using Algorithm 1 with $\alpha = 0$. The termination criterion suggested in Section VI was used with the norm chosen to be the supremum norm and a tolerance of 10^{-6}. Convergence was reached after three iterations.

Fig. 3 depicts the true least favorable densities and the ones estimated by means of the band model, the wobbly lines corresponding to the latter. The deviations of the estimates

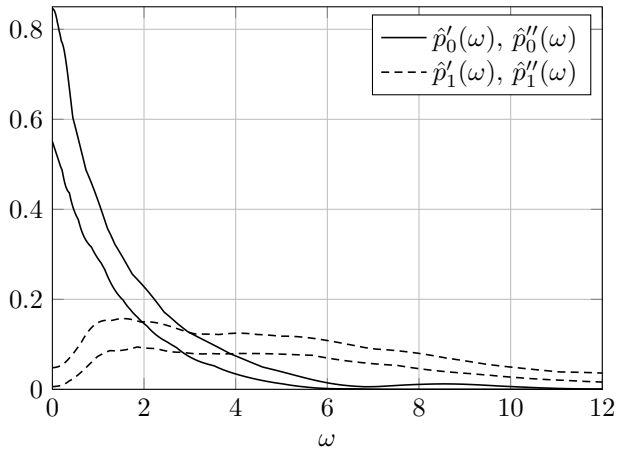

Fig. 2. Density bands estimated under $\mathcal{H}_0$ and $\mathcal{H}_1$, respectively, using 400 samples and 500 bootstrap instances of a gamma-kernel density estimator.

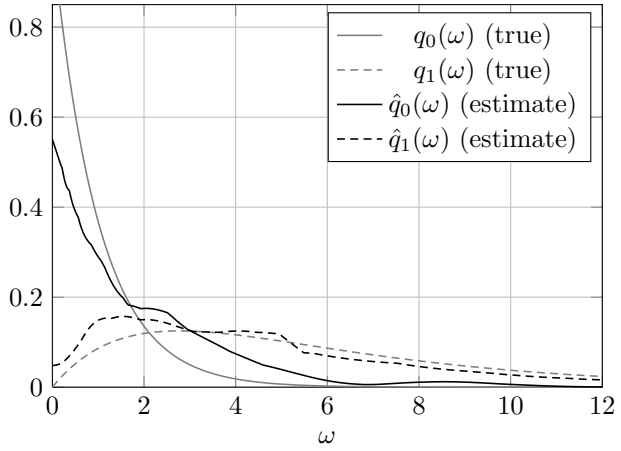

Fig. 3. True least favorable densities for the model given in (25) and (26) and estimates for least favorable densities based on the band model.

from the exact solution is clearly visible. However, given that a rather straightforward estimation approach and no *a priori* knowledge about the underlying distributions was used, the resemblance is reasonably close, especially under $\mathcal{H}_1$.

The same holds for the log-likelihood ratio depicted in Fig. 3. The estimated robust test statistic is more conservative, but approximately follows the optimal shape. The intervals of constant likelihood ratio stretch from around 1.5 to 2.5 and 5 to 5.5. The first one is located around the zero crossing of the optimal statistic and results in a censoring of observations with moderate power that cannot be reliably associated with a hypothesis. The second one causes an additional reduction of the influence of observations with very high power, making the test less sensitive to this kind of outliers.

The example demonstrates how the band model can be used in a purely data-driven manner to obtain robust test statistics that adapt to the underlying distributions and at the same time have a well-defined optimality property. We believe that our results help to better understand the resulting test and simplify its implementation by offering an efficient and generic algorithm to calculate least favorable densities for any numerically specified density bands.

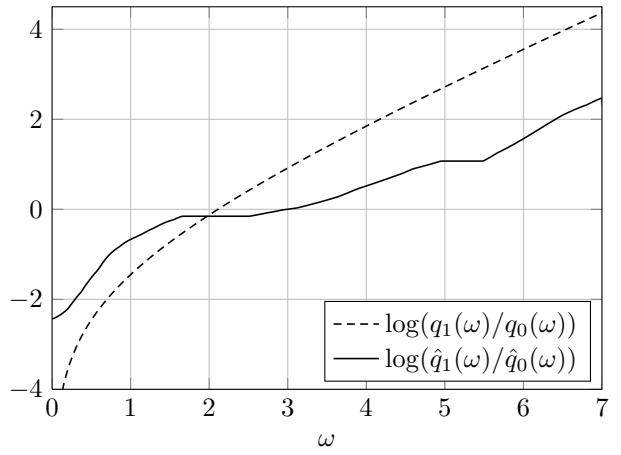

Fig. 4. True least favorable densities for the model given in (25) and (26) and estimates for least favorable densities based on the band model.

APPENDIX A
PROOF OF THEOREM 3

The proof of Theorem 3 is given in two steps. First, we show that (5) in Theorem 2 defines, up to some normalization constant, an f-divergence between P_0 and P_1. Second, it is shown that a pair of distributions that satisfies Theorem 3 minimizes all f-divergences over $\mathcal{P}_0 \times \mathcal{P}_1$ and hence satisfies Theorem 2.

An f-divergence $D_f(P_0||P_1)$ between two distributions P_0 and P_1 is defined via a convex function $f : \mathbb{R}_+ \to \mathbb{R}$ with $f(1) = 0$ and

$$D_f(P_0||P_1) = \int_{\{p_1>0\}} f\left(\frac{p_0}{p_1}\right) \mathrm{d}P_1 + f'(\infty)P_0[p_1 = 0], \tag{29}$$

where

$$f'(\infty) := \lim_{s\to\infty} \frac{f(s)}{s}.$$

The second term in (29) is often omitted under the assumption that P_1 dominates P_0. In [29], it is shown that every f-divergence can equivalently be defined via a convex function $\phi : [0, 1] \to \mathbb{R}$ with $\phi(0.5) = 0$ and

$$D_\phi(P_0||P_1) = \int \phi\left(\frac{p_0}{p_0 + p_1}\right) d(P_0 + P_1). \tag{30}$$

Moreover, given either ϕ or f, the complementing function is uniquely determined by

$$f(s) = (s + 1)\,\phi\left(\frac{1}{s+1}\right),\ s \in \mathbb{R}_+ \quad \text{and}$$
$$\phi(s) = s\, f\left(\frac{1-s}{s}\right),\ s \in [0, 1], \tag{31}$$

respectively. Therefore, with $\phi(s) = \Psi(s) - \Psi(0.5)$,

$$H_\Psi(P_0, P_1) = D_\phi(P_0||P_1) + \Psi(0.5)$$

defines an f-divergence with constant offset $\Psi(0.5)$, which is independent of P_0 and P_1.

The last step to prove Theorem 3 is to note that every f-divergence can be written as [30]

$$D_f(P_0||P_1) = c_f - \int L_\lambda(P_0, P_1)\, \mathrm{d}\nu_f(\lambda),$$

where c_f is a constant and ν_f is a nonnegative measure on $[0,\infty)$, both depending only on f. Therefore, if (Q_0, Q_1) maximizes L_λ over $\mathcal{P}_0 \times \mathcal{P}_1$ for all $\lambda \geq 0$, it minimizes at the same time all f-divergences and in turn satisfies Theorem 2.

Appendix B
Proof of Theorem 5

The proof of Theorem 5 is a straightforward application of Lagrangian duality. The optimization problem at hand is

$$\max_{p_0,p_1} \int \min\{p_0, \lambda p_1\}\,\mathrm{d}\mu \quad \text{s.t.} \quad \int p_i\,\mathrm{d}\mu = 1, \quad p_i' \leq p_i \leq p_i'',$$

for $i = 0, 1$. Replacing the minimum under the integral with two additional constraints yields

$$\max_{r,p_0,p_1} \int r\,\mathrm{d}\mu \quad \text{s.t.} \quad \int p_i\,\mathrm{d}\mu = 1, \quad p_i' \leq p_i \leq p_i'', \quad r \leq p_0,\ r \leq \lambda p_1.$$

The Lagrangian dual of this problem is

$$\min_{s,t,u,v} J(s,t,u,v) \quad \text{s.t.} \quad s,t,u \geq 0,$$

where

$$J(s,t,u,v) = \max_{r,p_0,p_1} \Big\{ \int r(1-s_0-s_1) + p_0(s_0+t_0-u_0-v_0) + \lambda p_1(s_1+t_1-u_1-v_1)\,\mathrm{d}\mu \Big\} + \int (w_0 + \lambda w_1)\,\mathrm{d}\mu + \lambda v_0 + v_1 \tag{32}$$

and we introduced the auxiliary functions

$$w_i = u_i p_i'' - t_i p_i', \quad i = 0, 1.$$

The real variables $v_i \in \mathbb{R}$ and the real-valued nonnegative functions $s_i, t_i, u_i \colon \Omega \to [0,\infty)$, $i = 0, 1$, are Lagrangian multipliers.[2] By strong duality,

$$L_\lambda(Q_0, Q_1) \leq J(s,t,u,v) \tag{33}$$

for all v and $s, t, u \geq 0$. Since J is unbounded unless the weighting functions associated with r and p under the integral in (32) are zero almost everywhere, the dual problem becomes

$$\min_{s,t,u,v} \int w_0 + \lambda w_1\,\mathrm{d}\mu + \lambda v_0 + v_1 \tag{34}$$
$$\text{s.t.} \quad s,t,u \geq 0, \quad s_0 + s_1 = 1, \quad s + t - u - v = 0.$$

Substituting $t = u + v - s \geq 0$ yields

$$w_i = u_i(p_i'' - p_i') + (s_i - v_i)p_0' \quad \text{with} \quad u \geq s - v.$$

[2]The multipliers are associated with the constraints as follows:

$$\underbrace{r \leq p_0}_{s_0}, \quad \underbrace{r \leq \lambda p_1}_{s_1}, \quad \underbrace{p_i' \leq p_i}_{t_i}, \quad \underbrace{p_i \leq p_i''}_{u_i}, \quad \underbrace{\int p_i\,\mathrm{d}\mu = 1}_{v_i}.$$

Apart from s_1, the multipliers associated with p_1 have been scaled by λ in (32), i.e., $v_1 \leftarrow \lambda v_1$, $u_1 \leftarrow \lambda u_1$, and $t_1 \leftarrow \lambda t_1$.

Since the objective (34) is non-decreasing in u, the last constraint holds with equality whenever $s - v$ is positive. We can hence write

$$u = \max\{s - v, 0\}$$

so that

$$\begin{aligned} w_i &= u_i(p_i'' - p_i') + (s_i - v_i)p_i' \\ &= \max\{s_i - v_i, 0\}(p_i'' - p_i') + (s_i - v_i)p_i' \\ &= \max\{s_i - v_i, 0\}p_i'' + \min\{s_i - v_i, 0\}p_i'. \end{aligned}$$

By (33), any feasible combination of dual variables provides an upper bound on L_λ. Assuming that

$$v_0, v_1 \in [0,1] \quad \text{and} \quad s_0, s_1 \in \{0,1\}$$

yields

$$w_i = \begin{cases} (1 - v_i)p_i'', & s_i = 1 \\ -v_i p_i', & s_i = 0 \end{cases}$$

so that

$$\begin{aligned} \lambda w_0 + w_1 &\geq \min\{(1-v_0)p_0'' - \lambda v_1 p_1',\ \lambda(1-v_1)p_1'' - v_0 p_0'\} \\ &= \min\{\hat{q}_0,\ \lambda\hat{q}_1\} - v_0 p_0' - \lambda v_1 p_1' \end{aligned} \tag{35}$$

with $\hat{q}_i$ defined in Theorem 5. Substituting (35) back into (34) yields the upper bound

$$\begin{aligned} L_\lambda(Q_0, Q_1) &\leq \int (\min\{\hat{q}_0,\ \lambda\hat{q}_1\}\,\mathrm{d}\mu - v_0 p_1' - \lambda v_1 p_1')\,\mathrm{d}\mu \\ &\quad + v_0 + \lambda v_1 \\ &= \int \min\{\hat{q}_0,\ \lambda\hat{q}_1\}\,\mathrm{d}\mu + \varepsilon_0 v_0 + \lambda\varepsilon_1 v_1. \end{aligned}$$

Appendix C
Proof of Lemma 1

We only detail the proof for g_0, the one for g_1 follows analogously. By inspection g_0 is nondecreasing in c_i. Moreover, since p_0 and p_1 both integrate to one,

$$g_0(c + \Delta c\,;\, \alpha p_0 + p_1) \leq g_i(c\,;\, \alpha p_0 + p_1) + (1 + \alpha)\Delta c$$

for all $\Delta c \geq 0$ so that g_i is continuous for all finite α. Finally,

$$g_0(0\,;\, \alpha p_0 + p_1) = \int p_i'\,\mathrm{d}\mu - 1 < 0$$

and

$$g_0(\tfrac{1}{\alpha}\,;\, \alpha p_0 + p_1) \geq \int p_0\,\mathrm{d}\mu - 1 = 0$$

so that

$$g_i(c\,;\alpha p_0 + p_1) = 0$$

for some $0 < c \leq \frac{1}{\alpha}$.

APPENDIX D
PROOF OF CONVERGENCE OF ALGORITHM 1

The following corollary will be useful for the proof.

Corollary 1: Let $P \in \mathcal{M}_\mu$. If a distribution Q^* with density

$$q^* = \min\{q'' \, , \, \max\{c\, p \, , \, q'\}\}$$

exists, it is unique and jointly minimizes

$$D_f(P\|Q) \quad \text{and} \quad D_f(Q\|P)$$

among all $Q \in \mathcal{Q}^= = \{Q \in \mathcal{M}_\mu : q'' \leq q \leq q'\}$ and for all convex functions f.

Corollary 1 follows from Theorem 4 by assuming either P_0 or P_1 to be fixed, i.e., $p' = p = p''$, and choosing $\alpha = 0$. Q^* is unique since band uncertainty sets are compact and $D_f(P\|Q)$ and $D_f(Q\|P)$ can be chosen to be strictly convex in Q.

Let us first consider the case when $\alpha = 0$ is a feasible choice. It then follows from Corollary 1 that for every f-divergence

$$\begin{aligned} D_f(Q_0^n\|Q_1^n) &= \min_{P\in\mathcal{P}_1^=} D_f(Q_0^n\|P) \\ &\leq D_f(Q_0^n\|Q_1^{n-1}) \\ &= \min_{P\in\mathcal{P}_0^=} D_f(P\|Q_1^{n-1}) \\ &\leq D_f(Q_0^{n-1}\|Q_1^{n-1}) \end{aligned} \tag{36}$$

so that $D_f(Q_0^n\|Q_1^n)$ is a monotonically non-increasing sequence of nonnegative real numbers that is guaranteed to converge by the monotone convergence theorem [31], i.e,

$$D_f(Q_0^{n-1}\|Q_1^{n-1}) = D_f(Q_0^n\|Q_1^{n-1}) = D_f(Q_0^n\|Q_1^n)$$

for $n \to \infty$. Since Q_0^n and Q_1^n are unique minimizers of $D_f(Q_0^n\|Q_1^{n-1})$ and $D_f(Q_0^n\|Q_1^n)$, respectively, this implies that in the limit $Q_0^{n-1} = Q_0^n$ and $Q_1^{n-1} = Q_1^n$.

For $\alpha > 0$, the proof is similar. Instead of general f-divergences, we consider the total variation distance, which is obtained by choosing $f(s) = f_{\text{TV}}(s) = |1-s|$ and can be written as

$$D_{f_{\text{TV}}}(P\|Q) = \int |p-q| \, \mathrm{d}\mu = \|p-q\|,$$

where $\|\cdot\|$ denotes the L^1 norm. Note that the total variation distance is symmetric and satisfies the triangle inequality. With $\pi := \frac{\alpha}{1+\alpha} \in (0,1)$, it is the case that

$$\begin{aligned} &\|q_0^n - q_1^n\| \\ &\overset{\text{(I)}}{\leq} \|q_0^n - \pi q_0^n - (1-\pi)q_1^{n-1}\| + \|q_1^n - \pi q_0^n - (1-\pi)q_1^{n-1}\| \\ &\overset{\text{(II)}}{\leq} (1-\pi)\|q_0^n - q_1^{n-1}\| + \|q_1^{n-1} - \pi q_0^n - (1-\pi)q_1^{n-1}\| \\ &= (1-\pi)\|q_0^n - q_1^{n-1}\| + \pi\|q_0^n - q_1^{n-1}\| \\ &= \|q_0^n - q_1^{n-1}\| \\ &\overset{\text{(III)}}{\leq} \|q_0^n - \pi q_0^{n-1} - (1-\pi)q_1^{n-1}\| + \pi\|q_0^{n-1} - q_1^{n-1}\| \\ &\overset{\text{(IV)}}{\leq} \|q_0^{n-1} - \pi q_0^{n-1} - (1-\pi)q_1^{n-1}\| + \pi\|q_0^{n-1} - q_1^{n-1}\| \\ &= \|q_0^{n-1} - q_1^{n-1}\|, \end{aligned}$$

where (II) and (IV) follow from Corollary 1 by construction of q_1^n and q_0^n, respectively. Hence,

$$\|q_0^{n-1} - q_1^{n-1}\| = \|q_0^n - q_1^{n-1}\| = \|q_0^n - q_1^n\|,$$

for $n \to \infty$. Finally, the triangle inequalities (I) and (III) only hold with equality if $Q_1^{n-1} = Q_1^n$ and $Q_0^{n-1} = Q_0^n$, which concludes the proof.

REFERENCES

[1] S. Kassam, "Robust hypothesis testing for bounded classes of probability densities (corresp.)," *IEEE Trans. Inf. Theory*, vol. 27, no. 2, pp. 242–247, 1981.

[2] P. J. Huber, "Robust estimation of a location parameter," *The Annals of Mathematical Statistics*, vol. 35, no. 1, pp. 73–101, 1964.

[3] ——, "A robust version of the probability ratio test," *The Annals of Mathematical Statistics*, vol. 36, no. 6, pp. 1753–1758, 1965.

[4] ——, *Robust Statistics*. Hoboken, New Jersey, USA: Wiley, 1981.

[5] S. Al-Sayed, A. M. Zoubir, and A. H. Sayed, "An optimal error nonlinearity for robust adaptation against impulsive noise," in *Proc. of the 14th IEEE Int. Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 2013, pp. 415–419.

[6] O. Zeitouni, J. Ziv, and N. Merhav, "When is the generalized likelihood ratio test optimal?" *IEEE Transactions on Information Theory*, vol. 38, pp. 1597–1602, 1992.

[7] B. C. Levy, "Robust hypothesis testing with a relative entropy tolerance," *IEEE Transactions on Information Theory*, vol. 55, pp. 413–421, 2009.

[8] G. Gul and A. Zoubir, "Robust hypothesis testing for modeling errors," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2013, pp. 5514–5518.

[9] E. Nikolaidis, Z. P. Mourelatos, and V. Pandey, *Design Decisions under Uncertainty with Limited Information*, 1st ed., ser. Structures and Infrastructures. Boca Raton, Florida, USA: CRC Press, 2011.

[10] F. Österreicher, "On the construction of least favourable pairs of distributions," *Zeitschrift fr Wahrscheinlichkeitstheorie und Verwandte Gebiete*, vol. 43, no. 1, pp. 49–55, 1978.

[11] B. C. Levy, *Principles of Signal Detection and Parameter Estimation*, 1st ed. New York CIty, New York, USA: Springer, 2008.

[12] S. Kassam and H. Poor, "Robust techniques for signal processing: A survey," *Proceedings of the IEEE*, vol. 73, no. 3, pp. 433–481, 1985.

[13] H. Poor, "Robust matched filters," *IEEE Trans. Inf. Theory*, vol. 29, no. 5, pp. 677–687, Sep 1983.

[14] R. Maronna, D. Martin, and V. Yohai, *Robust Statistics: Theory and Methods*. Hoboken, New Jersey, USA: Wiley, 2006.

[15] E. L. Lehmann and J. P. Romano, *Testing Statistical Hypotheses*, 3rd ed. New York City, New York, USA: Springer, 2005.

[16] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, 3rd ed., ser. McGraw-Hill Series in Electrical Engineering. New York City, New York, USA: McGraw-Hill, 1991.

[17] M. Sion, "On general minimax theorems," *Pacific Journal of Mathematics*, vol. 8, no. 1, pp. 171–176, 1958.

[18] D.-Z. Du and P. M. Pardalos, Eds., *Minimax and Applications*. Springer Science & Business Media, 1995, ch. 1, pp. 1–23.

[19] K. Seung-Jean and S. Boyd, "A minimax theorem with applications to machine learning, signal processing, and finance," in *46th IEEE Conference on Decision and Control*, 2007, pp. 751–758.

[20] H. Komiya, "Elementary proof for Sion's minimax theorem," *Kodai Math. J.*, vol. 11, pp. 5–7, 1988.

[21] V. Veeravalli, T. Basar, and H. V. Poor, "Minimax robust decentralized detection," *IEEE Transactions on Information Theory*, vol. 40, pp. 35–40, 1994.

[22] P. J. Huber and V. Strassen, "Minimax tests and the Neyman-Pearson lemma for capacities," *The Annals of Statistics*, vol. 1, no. 2, pp. 251–263, 1973.

[23] A. G. Dabak and D. H. Johnson, "Geometrically based robust detection," in *In Proc. of the Information Sciences and Systems Conference, (Baltimore, MD)*, 1993.

[24] T. Yucek and H. Arslan, "A survey of spectrum sensing algorithms for cognitive radio applications," *IEEE Commun. Surveys Tuts.*, vol. 11, no. 1, pp. 116–130, 2009.

[25] J. Mitola and J. Maguire, G. Q., "Cognitive radio: making software radios more personal," *IEEE Pers. Commun.*, vol. 6, no. 4, pp. 13–18, 1999.

[26] B. Shen, L. Huang, C. Zhao, Z. Zhou, and K. Kwak, "Energy detection based spectrum sensing for cognitive radios in noise of uncertain power," in *Proc. International Symposium on Communications and Information Technologies (ISCIT)*, Vientiane, Laos, Oct 2008, pp. 628–633.
[27] S. Chen, "Probability density function estimation using gamma kernels," *Annals of the Institute of Statistical Mathematics*, vol. 52, no. 3, pp. 471–480, 2000.
[28] A. M. Zoubir and D. R. Iskander, *Bootstrap Techniques for Signal Processing*. Cambridge, United Kingdom: Cambridge University Press, 2007.
[29] F. Osterreicher and I. Vajda, "Statistical information and discrimination," *IEEE Transactions on Information Theory*, vol. 39, no. 3, pp. 1036–1039, 1993.
[30] A. Guntuboyina, S. Saha, and G. Schiebinger, "Sharp inequalities for f -divergences," *IEEE Transactions on Information Theory*, vol. 60, no. 1, pp. 104–121, 2014.
[31] W. Rudin, *Real and Complex Analysis*, 3rd ed. New York City, New York, USA: McGraw-Hill, 1987.